

Characterising Fitness Landscapes Through Evolvability

Tom Smith^{*1,2}, Phil Husbands^{1,3} and Michael O'Shea^{1,2}

¹Centre for Computational Neuroscience and Robotics (CCNR)

²School of Biological Sciences

³School of Cognitive and Computing Sciences

The paper proceeds as follows: Section 2 outlines the concepts of fitness landscapes and neutrality, and goes on to describe the notion of problem difficulty as determined by the structure of the fitness landscape. Section 3 introduces the notion of solution evolvability as defined by local characteristics of the fitness landscape surrounding the solution, and derives and applies the problem difficulty metrics used in the remainder of the paper. Section 4 describes the tunably rugged and tunably neutral terraced NK landscapes used as test problems in this work. Sections 5 and 6 apply the metrics derived in section 3 to the test landscapes, and show that they predict the difficulty of searching in the tunably rugged and tunably neutral landscapes. Finally, section 7 shows that the metrics are robust when applied to online samples collected during simple hill-climbing, and the paper closes with discussion.

2 Fitness landscapes and neutrality

This section introduces two of the main concepts used in the paper. The *fitness landscape* (section 2.1), first introduced by Wright (1932), describes the search space as a multi-dimensional landscape defined by the genotype-to-fitness mapping through which evolution moves. The classical idea of searching this landscape for good genotypes focuses on the difficulty of climbing up to the globally optimal fitness solution, and avoiding locally optimal solutions. Here we argue that in difficult search problems, much of the time spent optimising may be spent in non-adaptive *neutral* evolution (section 2.2). Thus measures aimed at predicting the difficulty of search, must take account of the nature of neutrality in the space. Section 2.3 describes how the difficulty of finding good solutions is determined by the structure of the fitness landscape, and section 2.4 outlines different methods for sampling the fitness landscape structure.

2.1 Fitness landscapes

Wright (1932) introduced the *fitness landscape*





Figure 1: A two-dimensional model fitness landscape, with one globally-optimal and one locally-optimal peak. From a starting point, typically chosen at random, the search process tries to find good solutions. The process typically creates a new set of solutions through the application of genetic operators to the current solution(s), evaluating whether the new set is better than the current solutions. Evolving populations will tend to get stuck at the locally-optimal peak due to its large basin of attraction, and from there will only find the global optimum with difficulty.

A more exact picture, especially when dealing with solutions represented by discrete-valued genotypes, is the connected graph (Stadler, 1996). Solution vertices, or nodes, are connected directly through the action of the genetic operators. The graph may show the space in a very different way to the fitness landscape: mutation operators acting on more than one loci, and other operators such as recombination, may not 'see' fitness landscape hill-tops as local optima at all. However, local optima can clearly exist in the graph, occurring as graph nodes from which all connected nodes are of lower fitness. This definition may produce local optima with respect to genetic operators other than simply mutation, for example some solutions may be local optima with respect to recombination operators.


The graph definition of the search space highlights the dangers in the simple visualisable picture afforded to us by the fitness landscape description: our intuitive view may not apply in higher dimensional spaces. Fisher, for example, argued that local optima may not exist in a large class of high-dimensional spaces; the probability that a solution is optimal in every single dimension simultaneously is negligible (quoted in Provine, 1986, p. 274). However, it should be stressed that many problems clearly do show local optimality, e.g. the travelling salesman problem (Lawler et al., 1985). The next section introduces the idea of search space *neutrality*, one possible way in which some high-dimension spaces may differ radically from our intuitive viewpoint.

2.2 Fitness landscape neutrality


In the neutral theory, it is argued that evolving populations may spend relatively large periods of time



(a) Unconnected peaks



(b) Single neutral pathway



(c) Broad neutral plateau

Figure 2: Three two-dimensional model fitness landscapes showing the possible advantage of neutrality in a simple landscape with one globally-optimal and one (nearly) locally-optimal peak. shows the two peaks as unconnected; populations evolving to the locally optimal peak will have difficulty moving to the global optimum. shows the two peaks connected by a single neutral pathway; a population on the sub-optimal peak will eventually find the pathway. shows the two peaks connected by a broad plateau; the population will move easily from the sub-optimal peak to the global optimum.

secondary structure folding algorithms show that neutral walks (a neutral variant on the random walk,

3.1 The transmission function

$$\mathbf{E}_b = \int_{-\infty}^{\infty} \mathbf{f} T(\mathbf{f} : \mathbf{h}, \mathbf{k}) d\mathbf{f} \quad (5)$$

$$\text{where } F_c \text{ defined by } |G_{F_c}^+(\mathbf{h}, \mathbf{k})| = \frac{C |G(\mathbf{h}, \mathbf{k})|}{100} \quad (13)$$

The mean fitness of the set of o spring with fitness in the bottom percentile can be defined through the set $G_{F_d}^-(\mathbf{h}, \mathbf{k})$ of o spring with fitness below some fitness F_d .

The next section applies the metrics to a set of simple cases, where the parent genotypes lie at different points in a hypothetical landscape.

3.4 Simple evolvability examples

The metrics derived in the previous two sections are here applied to a set of simple cases, showing their

landscape increases, up to the maximally rugged random $K = N - 1$ landscape, corresponding to the random energy spin-glass model (Derrida, 1981). Early work by Weinberger (1990, 1991) has shown that increasing ruggedness of the system is well predicted by the decrease in correlation length for the system, and that the number of locally optimal peaks increases dramatically with K ; this has long been the benchmark result arguing that landscape ruggedness is the key feature for problem difficulty. Figure 6 shows the correlation lengths derived from random walks in the $N = 25$, $K = [0, 1, 2, 6, 12, 18, 24]$ landscapes; increasing K clearly produces more rugged landscapes with shorter correlation lengths.




Figure 6: Correlation lengths calculated over random walks.






Figure 7: Correlation lengths calculated over random walks on the neutrally terraced N -landscapes (mean and standard deviation over 1000 walks shown). Shown for $N = 25$, $F = 4$. The correlation lengths are identical to the corresponding lengths shown in figure 6 (calculated on the non-neutral N -landscape), and decrease with increasing K , showing

with mean and deviation dependent on N , K and the current solution fitness (see e.g. Weinberger, 1990; Stadler and Schnabl, 1992). From this it is possible to derive the expected fitnesses (and the time taken on both adaptive and random walks) at which local optima are reached for various N and K (again, see e.g. Weinberger, 1990; Stadler and Schnabl, 1992). In the next section we derive analytic and empirical results for the evolvability measures when applied to the NK landscapes.


5.1 Analytically derived evolvability for NK landscapes

In this section, we focus on the probability that an offspring derived from a single bit mutation of the parent has a higher (or equal) fitness than the parent, i.e. the first evolvability metr


erfc(




(a) Probability of a non-deleterious mutation, E




(a) Probability of a non-deleterious mutation, E



(b) Expected fitness over all mutations, E



significantly smaller than the probability of reaching a local optimum in the non-neutral $F = 0.5$. Rather than sticking in local optima, the search process can explore more of the space along neutral networks, eventually reaching higher fitness solutions.



7 Online sampling evolvability

In the previous sections we have investigated empirically derived evolvabilities for the tunably rugged and tunably neutral terraced NK landscapes through random sampling of the space of all solutions. This random sampling technique works well with the NK landscapes where solution fitnesses are defined as the linear sum of all loci fitnesses; due to the central limit theorem, the solution fitnesses will be normally distributed. However, in many problems, such normally distributed solution fitnesses will not be encountered, and metrics based on random sampling of the space will in general be less successful at predicting problem difficulty (see e.g. Smith et al., 2001a).

With such skewed solution fitness distributions, it may be necessary to bias the collected sample through only keeping a percentage of solutions found at each fitness, and define the problem difficulty metric over this biased sample. With even more extremely skewed distributions, it may be necessary to collect a biased sample through some direct search optimisation procedure such as a simple hill-climber. For example, Smith et al. (2001a) find only 0.0001% of random solutions have fitness above 50% of the

